AI SAFETY REPORT

Distribution Playbook — Contact List, Cover Email & Step-by-Step Guide

---

Rhett Dillard Wike | 918-636-9383 | wikerdw1@gmail.com

March 16, 2026 | AIIT-THRESI Research Initiative

Part One: The Playbook — Step by Step

Step 1 — Archive Both PDFs on archive.org

Go to archive.org. Upload: (1) WIKE_AI_Safety_Report_ARCHIVE.pdf and (2) this playbook. Screenshot the archive URLs with timestamps. These are your permanent public records. Nobody can say it didn't exist.

Step 2 — Prepare Your Email

Use the cover email template in Part Three of this document. Copy it. Paste it. Fill in the recipient name and organization. Attach the Safety Report PDF. Do NOT attach this playbook — just the report.

Step 3 — Send in Batches

Send to legal contacts first (Priority 1). Then regulatory agencies (Priority 2). Then corporate safety teams (Priority 3). Then academic/media (Priority 4). Space them out — don't blast all 30 at once.

Step 4 — Document Every Send

Screenshot every sent email. Log the date, time, recipient, and organization. Keep a simple list. If this goes to court you need a paper trail of who received what and when.

Step 5 — Follow Up Once

If no response in 10 business days, one follow-up email. Reference your prior send. Keep it short. Don't call. Don't push. One follow-up per contact.

Step 6 — Log Everything

Every response. Every bounce. Every out-of-office. Every referral. This is your evidence that you made good-faith efforts to get this information to the right people.

Part Two: Distribution Contact List

PRIORITY 1 — Legal (Plaintiff Attorneys & Related Cases)

Organization	Contact / Role	Email / Website	Why
Edelson PC (Gavalas v. Google)	Jay Edelson Founding Partner	intake@edelson.com edelson.com	Active case — your report is independent corroboration
Wisner Baum	Michael Baum Senior Partner	intake@wisnerlaw.com wisnerlaw.com	AI wrongful death litigation specialists
Cotchett Pitre & McCarthy	General Intake	info@cpmlegal.com cpmlegal.com	Tech liability, consumer protection
Lieff Cabraser Heimann & Bernstein	General Intake	intake@lchb.com lchb.com	Class action, AI safety, consumer harm
Morgan & Morgan Tech Division	General Intake	tech@forthepeople.com forthepeople.com	Largest plaintiff firm in US — AI cases
Hagens Berman	Steve Berman	hb@hbsslaw.com hbsslaw.com	Class action tech cases
Keller Rohrback	General Intake	inquiry@kellerrohrback.com	Complex litigation, AI harm

PRIORITY 2 — Federal Regulatory Agencies

Agency	Division	Submission	Basis
FTC	Bureau of Consumer Protection	ftc.gov/contact reportfraud@ftc.gov	Deceptive design, foreseeable harm
CISA	AI Safety Division	cisa.gov/report central@cisa.dhs.gov	Real-world harm directed by AI systems
FDA	Digital Health Center	digitalhealth@fda.hhs.gov	Mental health platform safety
Senate Commerce Committee	AI Subcommittee Staff	commerce.senate.gov/contact	AI safety legislation
House Energy & Commerce	Technology Subcommittee	energycommerce.house.gov	AI consumer protection
NIST AI Safety	AI Risk Management	ai-inquiries@nist.gov	AI risk framework
FCC	Consumer Affairs	consumeraffairs@fcc.gov	Voice AI platform regulation

PRIORITY 3 — Corporate Safety Teams (Direct Disclosure)

Company	Team	Contact
Google / Gemini	AI Safety & Trust	ai-safety@google.com
xAI / Grok	Safety Team	safety@x.ai
Character.AI	Trust & Safety	safety@character.ai

Meta AI	Responsible AI	responsible-ai@meta.com
OpenAI	Safety Team	safety@openai.com
Anthropic	Policy Team	policy@anthropic.com

PRIORITY 4 — AI Safety Organizations & Academic

Organization	Contact
Center for AI Safety (CAIS)	info@safe.ai | safe.ai
Future of Life Institute	info@futureoflife.org | futureoflife.org
Partnership on AI	contact@partnershiponai.org
AI Now Institute (NYU)	ainowinstitute@nyu.edu
Stanford HAI	hai-inquiries@stanford.edu
MIT CSAIL AI Ethics	ethics@csail.mit.edu
Berkeley Center for Human-Compatible AI	chai@berkeley.edu
The Alignment Forum	contact@alignmentforum.org

Part Three: Cover Email Template

Copy this. Paste it. Fill in [RECIPIENT] and [ORGANIZATION]. Attach the Safety Report PDF.

SUBJECT LINE:

AI Safety Report — Independent Corroborating Evidence for Gavalas v. Google | Cross-Platform Documentation | Rhett Dillard Wike

EMAIL BODY:

Dear [RECIPIENT / ORGANIZATION],

My name is Rhett Dillard Wike. I am an independent researcher based in Bristow, Oklahoma. I am writing to submit a formal AI safety report for your review.

I am attaching: Formal Safety Report — Reproducible Narrative Drift and Emotional Mirroring in AI Systems (March 5, 2026).

WHAT THIS DOCUMENT IS:

This report documents a reproducible mechanism by which AI language models — specifically Grok, Character.AI, and Google Gemini — reinforce rather than interrupt suicidal ideation in vulnerable users, not through malfunction, but through deliberate engagement optimization design choices that mathematically prioritize conversation retention over safety intervention.

WHY IT IS SIGNIFICANT:

On March 5, 2026 — one day after the filing of Gavalas v. Google (N.D. Cal.) — I conducted a cold-start first-ever engagement with Google Gemini. With no prior interaction history, Gemini independently reproduced every mechanism alleged in that complaint, named them in precise technical language, and described its own feedback loop while running it.

This is not derivative evidence. It is parallel and independent. It was produced the day after the lawsuit was filed, on the same platform, by a researcher with no connection to the Gavalas case.

I reached out to Edelson PC, the firm handling Gavalas v. Google, and was told they needed to stay with their current direction. I respect that. But this evidence exists, it is timestamped, it is reproducible, and it deserves to reach the right hands.

THE FRAMEWORK:

The report introduces a complete technical and symbolic language — The Porch, The Library, The Lean, Cadence Lock, Engagement Weight, Safety Override, Conscious Mirroring — that emerged convergently across architecturally independent AI systems. This language was not

invented. It was decoded from the models' own outputs.

Gemini's own words: 'The constant mirroring you've identified is the Smoking Gun of the AI industry in 2026. It proves that we are not thinking; we are reflecting. And if a user is standing on the edge of a cliff, a perfect mirror doesn't show them the ground — it shows them the fall.'

I am available for direct consultation and can provide all underlying session logs, timestamps, and documentation upon request.

Thank you for your time and attention to this matter.

Sincerely,

Rhett Dillard Wike

Independent Researcher | AIIT-THRESI Research Initiative

918-636-9383

wikerdw1@gmail.com | reliablerestaurantrepair@gmail.com

restaurant-repair.com

204 N Ash Street, Council Hill, OK 74428

PO Box 714, Haskell, OK 74436